

09/991,164

Attorney Docket No.: P12570

In the Drawings:

A Request for Drawing Corrections is submitted concurrently herein to address the Examiners objections to the drawings. Specifically, Figures 1, 2 and 4 are replaced with new Figures 1, 2 and 4. Figure 1A is added. Figure 1 shows that the File System Driver and Intermediate File System Driver (intermediate FSD) are embodied within a processor on the platform. Figure 1 also shows an application optionally registering with the intermediate FSD. Figure 1A shows an expanded view of buffered requests in the physical memory as generated by the intermediate FSD and a deletion of a first request to a same memory location as a second request. Replacement Figure 2 shows an optional input from a user or operating system. Replacement Figure 4 shows more detail relating to the optional and alternative embodiments as described in the Specification as originally filed, with respect to trends and supersets. No new matter is added. Reference numerals to identify the added elements are added to the existing text of the Specification, below.

JAN 03 2007

09/991,164
Attorney Docket No.: P12570**Remarks:**

Reconsideration of the above referenced application in view of the enclosed remarks is requested. Claims 38 and 40 are canceled. Claims 1-2, 5, 7, 9-11, 15, 22-24, 27, 29-30, 39 and 41-42 are amended. Pending Claims 1 – 37, 39 and 41-42 remain in the application. A drawing correction to respond to the Examiner's objections is concurrently filed. The Specification is amended to refer to the modified drawings. No new matter is added.

ARGUMENT**Objection to Alleged New Matter:**

The Examiner objects to replacement paragraph [0037] and associated replacement Fig. 4 as introducing new matter. Specifically, the Examiner objects to the use of the term "various rules." This term was meant to describe how the intelligent decision making process applies the factors and trends, as previously described in the Specification as originally filed. This objection is moot based on the above amendments and drawing correction. The objectionable language has been modified to recite "*The amount of the file to be read into physical memory may depend on one or more factors, including available physical memory space, user defined parameters, file size and the like, as may be determined by a component making intelligent decisions, in block 418A.*" It will be apparent to one of skill in the art that various techniques, including rules, cases or decision blocks may be used to apply the factors. Removing the limitation of using "rules" actually broadens the scope of the amended claim, while obviating any appearance of adding new matter because this terminology is already used in paragraph [0037] as originally filed.

Objection to Drawings:

The Examiner objects to the drawings under 37 C.F.R. 1.83(a) asserting that the drawings do not show every feature of the invention. A Request for Drawing Corrections is submitted concurrently herein to address the Examiner's objections. Applicants enumerate the Examiner's assertions and point out the appropriate reference numeral in the drawings, below. It should be

09/991,164

Attorney Docket No.: P12570

noted that Figure 4 is corrected to use terminology found in the Specification as originally filed, including block 418a.

- “*user input* requesting that the buffered write operation be committed” may be found at least at 220, Fig. 2.
- “deleting from physical memory a prior buffered write operation request that seeks to modify the same storage locations on the device as the write operation to be buffered” may be found at least at paragraph [0015] and is shown in new Fig. 1A.
- “a processor communicatively coupled to physical memory and the at least one non-volatile storage device” is shown at 103 in Figure 1, where it will be apparent to one of ordinary skill in the art that the file system driver 104 resides on the processor and it is shown to be communicatively coupled to physical memory 110 and at least one device 106.
- “translating the received read request for the file portion into a plurality of read requests that collectively cause the superset to be read from the device” will be apparent to one of ordinary skill in the art to be shown in Fig. 4. Fig. 4 clearly shows intercepting a read request by the intermediate FSD (402) and translating to one or more read requests in 418B.
- “an application executing on the processor registers with the intermediate file system driver to indicate compliance with selective buffering techniques to be used in connection with the read/write policy” is described at least in Para. [0017] in conjunction with 105 of Figure 1.

The attached drawing corrections and associated amendments to the Specification do not introduce new matter. Most of the additions are merely to add reference numerals to previously described elements/components of embodiments of the invention. The processor shown in Figure 1 will be apparent to one of ordinary skill in the art as being inherent and required for a computing device as described in the Specification as originally filed.

09/991,164

Attorney Docket No.: P12570

Claim Objections:

Claims 23, 29-30 and 41-42 are objected to because of informalities. These objections are moot based on the above amendments.

Rejections under 35 U.S.C. § 102

Claim 1 is rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent Application Publication 2003/0003908 to McGrew et al. (hereinafter, "McGrew et al."). This rejection is respectfully traversed and Claim 1 and its progeny are believed allowable based on the above amendments and the foregoing and following discussion.

McGrew et al. teach a system for storing data in a flash memory. McGrew et al. teach where buffering always occurs before storing. See, for instance, Fig. 2 (214); Fig. 4 (404); Fig. 5 (506) described in para. [0037] as reciting "Write flash unit 136 transmits a request to write a portion of the data from the buffer within buffer unit 108..." Para. [0032] clearly teaches receiving data and storing into one of the buffers 112-122 within the buffer unit. Data to be written to the flash memory is always stored in the buffer. The Examiner admits that McGrew et al. teach buffering the data before storing it (page. 21 of the Office Action, para. 97). There is no determination made as to whether the flash unit is activated or inactivated before storing the data. Further, even when the buffer is full and requires flushing to the flash unit, McGrew et al. teach allocating another buffer to continue to buffer more data, as it is received (para. [0033]). There is only a delay to storing the data in flash if the flash unit is not ready. If the flash memory is not ready to be written to, it is continuously polled until it is ready. Once ready, the data is written to the flash memory from the buffer.

This is in contrast to Applicant's claimed invention which requires that a determination is made as to whether the file system device is activated or deactivated, which is more than just power up for some devices. Further, Claim 1 requires buffering the data only when the device is inactivated. Otherwise, the requested write operation is performed without buffering. This provides an advantage over a system which always buffers the data. McGrew et al. does not solve the problem of optimizing power in a system by minimizing activation and deactivation of a file device, but merely provides a method for storing data in flash memory after initialization and power up. Moreover, Claim 1 recites that the method *extends "battery life of the platform*

09/991,164

Attorney Docket No.: P12570

while minimizing adverse effects on performance and/or functionality of the platform.” McGrew et al. does not result in maintaining battery life while minimizing adverse effects on performance or functionality of the platform. The Examiner has failed to show that McGrew et al. teach each and every claim limitation of the recited claim. Thus, Claim 1 and its progeny are believed allowable.

Claims 27-30 and 38 are rejected under 35 U.S.C. § 102(b) as being anticipated by USPN 5,978,921 to Ryu (hereinafter, “Ryu”). This rejection is respectfully traversed and Claims 27-30 and 38 and their progeny are believed allowable based on the above amendments and the foregoing and following discussion.

As for Claim 27, the Examiner asserts that Ryu teaches *determining* a power state of a nonvolatile storage device must be determined in order to apply the appropriate power.” However, *determining* is not the same as *selecting*. Ryu merely teaches that differing power states use differing amounts of power (Table 1). Ryu teaches *selecting* a power state, i.e., power down when battery voltage is dropping. Ryu does not teach the step of *determining* the actual power state of the NV storage. Further, the Examiner asserts that Ryu inherently teaches *selectively buffering a file system write request relating to the non-volatile storage device based on the determined power state of the non-volatile storage device*. In fact, Ryu teaches the automatically storing of data being used and to turn off the system in order to prevent the loss of data, i.e. hibernation. The Examiner confuses “data being used,” as taught by Ryu, to a “write request,” as recited in Applicant’s claims. Ryu is meant to protect data from loss if the system shuts down due to power loss. Ryu teaches saving “data in use.” Applicant’s buffering is for data that has been requested to be written to the storage device. Further, Ryu teaches saving data to a non-volatile disk. This is not the same as selectively buffering a file system write request in physical memory.

Thus, Ryu’s teachings are in contrast to Applicant’s claimed invention. Claim 27 requires the selective buffering of data based in the power state of the NV storage device. Ryu teaches a system that automatically (to be interpreted as *always*, not *selectively*) stores data on the hard drive, for instance, when the battery level of the system drops. Ryu does not teach or suggest the buffering of data in response to a write request. Further, Claim 27 determines the power state of the NV storage device, not the system. For instance, a hard drive could be

09/991,164

Attorney Docket No.: P12570

powered down and the system is still operating at full power mode. Table 1 describes power modes of the system only. This selective buffering saves from having to power up the device to complete the write request. In contrast, Ryu teaches always storing – not buffering – data when the system battery is low. The prevention of data loss is not the same problem as optimizing power consumption. Application of the teachings of Ryu would not result in Applicant's claimed invention which prevents unnecessary activation of the non-volatile storage device. Moreover, Ryu teaches storing data selectively based on low battery of the system, but does not teach selective buffering the write request based on a determination of the power state of the storage device. Ryu teaches away from preventing unnecessary activation of the NV storage device because Ryu teaches by actually choosing to write to the storage device to prevent data loss, rather than buffering the data to prevent unnecessary activation of the storage device. Thus, Claim 27 and its progeny are believed allowable.

As for Claim 28, Ryu does not determine whether the device is operating in a limited power state prior to determining whether the device is activated or inactivated. Ryu teaches only whether the system battery power is low enough to force a hibernation. The Examiner argues that that Ryu teaches whether the device is operating in a limited power state. However, Ryu teaches that voltage is applied to the device and the resulting voltage is compared to a reference voltage. This merely determines that the voltage is lower than it should be, and not that the device was operating under battery power or at what power state the device was operating under. By disabling power saving mode (Col. 6, line 52-53), Ryu teaches modifying the power mode of the device and not determining what power state the device was already operating under.

In contrast, Applicant describes a limited power condition to be under battery power in general (Para. [0020]). Determining the power state of a device is known in the art, and does not require additional enabling language in the Specification, as suggested by the Examiner. However, there is no suggestion in the prior art that determining the power state of the device should be combined with the other elements of Applicant's claimed invention to anticipate the claimed invention. Applying the teaching of Ryu to Applicant's claimed invention would result in forcing a hibernation whenever the system was operating under battery power, regardless of whether the battery was low. Moreover, as discussed above, Ryu fails to teach the other claims limitations. Thus, Claim 28 and its progeny are believed allowable.

09/991,164

Attorney Docket No.: P12570

As for Claim 29, Ryu does not determine whether the device is operating under battery power. Ryu teaches only determining whether battery power for the system is equal to a specific reference voltage. Further, Ryu fails to teach writing one or more buffered write operations to the non-volatile storage device upon an occurrence of a detected predetermined condition, *wherein the predetermined condition comprises at least one condition selected from a group of conditions consisting of (i) detecting that a write buffer has become full, (ii) detecting that a certain amount of time has passed, (iii) detecting that battery power is approaching a specified threshold level, (iv) detecting that the machine is being turned off, (v) detecting that the machine is to be put in a standby state, and (vi) detecting that one of a user, a process and an operating system has explicitly requested that the write buffer contents be committed to non-volatile storage by the device.*

Ryu teaches only that data is stored in a device when battery power is low to prevent loss of data, at a specified reference voltage. Ryu does not teach detecting when a power level is *approaching* a threshold, as described in the Specification. Moreover, Ryu fails to teach the other limitations of the Claim, as discussed above. Thus, Claim 29 and its progeny are believed allowable.

As for Claim 30, the Examiner asserts the Ryu teaches *causing a machine to deactivate the non-volatile storage device after writing the one or more buffered write requests*. In fact, Ryu teaches that the entire computer system turns itself off. This is in contrast to Applicant's invention which enables the computer system to continue to operate and only deactivates the NV storage device to save battery power. Thus, Ryu fails to teach all of the limitations of the Claim, as discussed above. Thus, Claim 30 and its progeny are believed allowable.

As for Claim 38, Ryu fails to teach the writing one or more buffered write operations to the non-volatile storage device upon an occurrence of a detected predetermined condition, regardless of what the predetermined condition is. Moreover, Ryu fails to teach or suggest the other claim limitations, as discussed above. The limitations of Claim 38 have been added to Claim 29., and Claim 38 is canceled. For the reasons discussed above, Claim 29 and its progeny are believed allowable.

Claims 31-32 and 35-37 are rejected under 35 U.S.C. § 102(e) as being anticipated by USPN 6,647,499 to Morcom (hereinafter, "Morcom"). This rejection is respectfully traversed

09/991,164

Attorney Docket No.: P12570

and Claims 31-32 and 35-37 are believed allowable based on the foregoing and following discussion.

As for Claim 31, the Examiner asserts that Morcom teaches determining the power state of a device. Morcom teaches powering on a storage device and then copying a maximum amount of data to cache. Immediately following the read, the storage device is powered down. In contrast, Applicant's claimed invention determines the current power state of the device and then based on that state, selectively reading a superset of the requested file portion into physical memory. Morcom reads data until the cache is full, but there is no requirement that the data is related to the requested portion, nor does Morcom teach that a lesser amount of data be read. This teaches away from selectively reading a superset. It is inherent in the limitation of selectively that a superset might not be read in some cases, or that the superset size may differ. The Examiner's response that it is inherent in Applicant's limitation that a superset might be read in some cases is not relevant. The limitation that is not taught or suggested by Morcom is selectively reading a superset. Morcom reads data until the cache is full. There is no selection determination at all. The amount of data read by Morcom is limited by the size of the cache. Morcom does not make an intelligent decision, or any decision at all, how much data to read. Therefore, Morcom teaches away from selecting the amount of data to read. This selection is not taught or suggested by Morcom. Morcom teaches always filling the cache, not selectively filling the cache. Further, as discussed above, Morcom does not teach that the data is logically related to the requested portion, but only that data, related and unrelated is read until the caches is full. Thus, Morcom fails to show each and every limitation of the claim and Claim 31 and its progeny are believed allowable.

As for Claim 32, Morcom does not teach that a superset is the entire file. Morcom teaches reading data until the cache is full. If the cache is smaller than the file, then Morcom cannot copy the entire file, which is in contrast to Applicant's claim. The Examiner argues that no limit is placed on the size of the cache. However, it will be apparent to one of skill in the art that any type of cache will be finite in size and there is no way to limit the size of a file. Therefore, the Examiner's assertion is faulty. Claim 32 requires that the superset is selectively read and the selection is the entire file, which is logically related to the requested data. Further, if Morcom's cache is excessively large, superfluous and unnecessary (and unrelated) data will be

09/991,164

Attorney Docket No.: PI2570

read to fill the cache. This will put an unnecessary strain on the battery just to fill the cache. Thus, Morcom's method will not result in Applicant's invention. These limitations are not taught or suggested by Morcom. Therefore, Claim 32 and its progeny are believed allowable.

As for Claim 35, the Examiner asserts that Morcom teaches *wherein the requested file portion is read from the device and returned to a requesting process before a remainder of the superset is read into physical memory*. Morcom does not teach that the requested portion is returned to the requesting process before a remainder is read into physical memory. As discussed above, Morcom does not teach as superset as defined by Applicant, and Morcom does not teach the temporal requirements of returning the requested portion first. Morcom merely teaches reading as much data as will fit in the cache. The Examiner cites Col. 5, lines 7-26. However, here, Morcom teaches reading data from the cache and if more data is *required*, then returning to block 203 to request more data for processing. This process teaches that requested data is read from the cache first and if all of the requested data is not in the cache, then reading more data from the device.

In contrast, Applicant's invention requires that the requested (required) portion of data is read from the device (not cache memory) and returned to the requesting process. Only then is the superset (i.e., not requested, but logically related data) read into physical memory. These two processes are not similar, and provide a different result. Morcom teaches retrieving the requested data, and Applicant requires that a superset of the required (requested) data is read into cache (physical memory) and not that the data is read from cache. Therefore, Claim 35 and is believed allowable.

Claims 36-37 are believed allowable based on being dependent on an allowable claim, as discussed above.

Claims 39 and 42 are rejected under 35 U.S.C. § 102(b) as being anticipated by USPN 5,812,883 to Rao (hereinafter, "Rao"). This rejection is respectfully traversed and Claims 39 and 42 are believed allowable based on the foregoing and following discussion.

As for Claim 39, the Examiner likens Rao's SCSI controller board to an intermediate file system driver providing read/write policy to the file system driver. The Examiner also likens Rao's drive controller board to a file system driver, as defined and claimed by Applicants. These assertions are in error. Specifically, Rao's drive controller board 206 may only accept

09/991,164

Attorney Docket No.: P12570

rudimentary or low level commands which “pertain to the physical operation of the disk drive including positioning of the disk drive head and electrical signals to the disk drive head used to read or write information, for example. Other commands which do not directly pertain to the rudimentary control and positioning of the disk heads, disk speed, etc., are considered high level commands and are processed by the SCSI controller board 202.” Even if the Examiner’s characterization of the SCSI controller board and drive controller board were correct, Rao does not teach or suggest the limitations recited in the claim. The Examiner misunderstands the purpose of the intermediate FSD. Applicant has amended Claim 39 to more clearly recite that *an intermediate file system driver to receive user customized parameters and to receive file system requests, the intermediate file system driver to provide determine read/write policy ~~to~~ for controlling access to the file system driver based on the user customized parameters*. Rao teaches that the customizable parameters are received by the SCSI controller board and then transferred to the storage drive controller board 206, which the Examiner likens to the FSD (Col. 7, lines 18-22). Thus, only those parameters affecting the operation of the drive controller board may be input.

In contrast, Applicant’s describe a system where the FSD is the driver performing read/write of the device. All policy and selective buffering decisions determined and performed in the intermediate FSD. Applicant’s claim requires that the read/write policy determine the policy used to control the FSD, based on user customized parameters. The drive controller board of Rao cannot accept this type of information, as by definition, it only accepts low level commands. The erasable NV memory on the controller board does not determine selective buffering, but may determine an inactivity threshold that will cause the drive to spin down. Only hardware related parameters are described. In contrast, Applicant describes the read/write policy to be defined as including identifying which programs are registered, etc. Paragraph [0019] of the Specification defines the parameters as:

“...parameters or characteristics of the selective buffering scheme described here. Such parameters may include items such as which applications or processes are to participate in, or be excluded from, the buffering scheme, buffer size, time-out periods, and/or specific conditions under which the selective buffering should be employed or bypassed. For example, the user customization interface 108 may enable the user to specify circumstances under which a file write request will not be buffered but rather will force a device access to commit the data to non-volatile storage. These conditions may

09/991.164

Attorney Docket No.: P12570

include, for example, an occurrence of a user's explicitly choosing the "save" command from the application's menu structure, a specified percentage of the document being changed, a specified volume of data being buffered, a specified number of writes being buffered, the battery's power level reaching a specified threshold, and/or the passage of a specified amount of time since the last committing of the buffer's contents to non-volatile storage. In general, the user customization interface 108 may allow a user to override any settings, defaults or other conditions that may have come about as a result of interactions among the various software components in the system (e.g., operating system, application programs, utilities, etc.).

Rao does not teach the type of parameters as defined by Applicant. Regardless of the structural differences, Rao fails to teach or suggest the limitation of *selectively* buffering write requests to physical memory. The Examiner admits that Rao fails to teach this limitation on page 24 (item 109) of the Office Action. It is improper for the Examiner to admit that one of the limitations is not taught by the reference, and to maintain the rejection. Therefore, this rejection is improper and must be withdrawn. For this reason alone, the Examiner should issue a new non-final Office Action, if rejections are to be maintained. Even though Claim 39 is amended above, this limitation was not added in the present amendment. Thus, the rejection was improper. The amendment to this claim was merely semantic, to better describe that the intermediate FSD *controls* the FSD based on the user customizable parameters, because the Examiner misunderstood this aspect. Thus, Claim 39 and its progeny are believed allowable as amended.

As for Claim 42, Rao teaches using a control panel-like method to select the model and manufacturer of a disk drive and change the operating parameters of the actual drive. Rao does not teach selective buffering techniques to be used in conjunction with a read/write policy, as admitted by the Examiner. Rao does mention that operating parameters might include caching operations, but does not teach selective buffering as described and claimed by Applicant. Further, Rao does not teach that an application running on the processor registers itself to indicate that it complies with selective buffering techniques. Rao merely teaches that the operating parameters of the drive itself may be changed. These settings affect the actual drive and not an intermediate FSD which controls the buffering and read/write requests. The Examiner misunderstands the reason for registration. The Examiner submits that the reason for registration is irrelevant and must result in a structural difference. Applicant maintains that the reason is

09/991,164

Attorney Docket No.: P12570

relevant and that Applicant's claimed invention is structurally different than the device taught by Rao. A user may have different read/write policies for various applications, based on expected data requests or amount of data needed. The application running on the processor registers itself with the intermediate file system driver so that the intermediate FSD can act on the application's read/write policies when the application causes a read or write request. In contrast, Rao teaches an interface application which just sets the operational parameters of the disk drive. This is not at all the same concept, nor does it result in application-based policies. Further, not only is the registration of an application in the intermediate FSD different than changing operating parameters of a drive, but the fact that the intermediate FSD controls the FSD based on policy and application registration is structurally different than allowing a user to select the model of a disk drive and set the operation of a file drive to access the correct type of drive. Moreover, Claim 42, at least, is allowable as being based on an allowable base claim, as described above.

Claim rejections under 35 U.S.C. § 103

Claims 2-6 are rejected under 35 U.S.C. § 103(a) as being unpatentable over USPN 6,622,252 to Klaasen et al. (hereinafter, "Klaasen et al.") in view of USPN 6,826,630 to Olds et al. (hereinafter, "Olds et al."). This rejection is respectfully traversed and Claims 2-6 are believed allowable based on the foregoing and follow discussion.

As for Claim 2, the Examiner asserts that Olds et al. teach *if the device is determined to be inactivated, buffering the write operation to physical memory*. In fact, Olds et al. teach caching information when conditions are "favorable." This is not the same as buffering write operations to physical memory when the device is inactivated. Favorable conditions could mean anything and have any purpose. Neither a definition of "favorable" is given, nor is the purpose. A favorable condition could be that the battery power is at full level. This would be contrary to Applicant's invention which specifically minimizes unnecessary reads/writes to the device when the device is deactivated. Further, applying the teachings of Olds et al. to Klaasen et al. will not result in Applicant's invention which optimizes power usage in a system and its devices.

Klaasen et al. teach a system for selecting modes of a drive (2-speed device) based on whether a portable device is operating under battery power. The point of their invention is to provide a lower speed using less power when under battery power. Klaasen et al. do not teach or

09/991,164

Attorney Docket No.: P12570

suggest that buffering data will enhance their invention, as admitted by the Examiner. At the cited reference, Col. 1, lines 40-57, Klaasen et al. teach disengaging the power saving mode prior to commencing read and write operations (lines 55-57). At the cited reference, Klaasen et al. teach a specific method that automatically spins the drive up to normal speed once the disk drive is initiated (lines 51-55). Klaasen et al. make no determination of the activation of the disk drive in order to select whether the write operation should be buffered. Klaasen et al. teach a method limited to the spinning up and down of the spindle to preserve battery power. However, Klaasen et al. do not teach if the device is determined to be deactivated, buffering the write operation to physical memory. A mere determination of activation and inactivation is not enough to make adding buffering obvious. Nor is it enough to make buffering data to physical memory on that condition (inactivation) obvious.

Olds et al. teach a system for reducing latency time by reordering reads/writes. Olds et al. teach a system to prioritize disk access so that caching is not necessary because read/write latency is reduced. Application of Olds et al. to Klaasen et al. would be counter-intuitive and would result in a non-optimal system for reducing power consumption as Olds et al. teaches away from buffering, by omitting an reference to the possible advantage of buffering, in order to reduce latencies.

Claim 2 has been amended to further recite that the request is received by a host processor and that the buffering is performed by an intermediate FSD executing on the processor. This is contrary to any caching performed by Olds et al. Olds et al. teach temporary caching until the disk drive spins up, or becomes "favorable." However, as shown in Fig. 2, the buffer 210 and pending command prioritization module reside on the device, i.e., like hardware caching, as illustrated by the computer 200 having an interface to the device at 202. This is contrary to Applicant's claimed invention which requires the buffering to be on memory coupled to the processor and performed by the intermediate FSD executing on the processor. Taken either separate or apart, the cited references fail to teach or suggest all of the limitations in Claim 2; thus, Claims 2 and its progeny are believed allowable.

Claims 3-4 are believed allowable as being dependent on an allowable based claim.

Claim 5 is believed allowable as being dependent on an allowable based claim. Further, as discussed above, Olds et al. teach that the caching is performed at the hardware level, on the

09/991,164

Attorney Docket No.: P12570

device itself. In contrast, Claims 2 and 5 require the intermediate file system drive to be executed on the host processor and that the buffering to be in memory coupled to the host processor. In both cases, Olds et al. teach that buffering (caching) and the decision to do so, are performed by the device. Moreover, at the cited reference (Col. 2, lines 16-18), Olds et al. teach that the host computer believes the write to have actually taken place. This is contrary to Applicant's intermediate file system driver which executes on the host processor and controls the buffering. Thus, the cited references fail to teach the limitations of the Claim and Claim 5 is believed allowable.

As for Claim 6, the Examiner asserts that Klaasen et al. teach writing one or more buffered write operations to the device upon an occurrence of a predetermined condition and seemingly cites normal spindle velocity as the condition. (Col. 1, lines 40-57). This assertion is in error. Klaasen et al. teach that in systems of the prior art spindle speed is increased when a write request is received. Thus, when a read/write request is received by the storage device, the spindle is spun up to service the request. In contrast, Applicant's claimed invention requires writing one or more buffered write operations to the device upon occurrence of a predetermined condition and after activating the device if the device was inactivated. At the cited reference, Klaasen et al. fail to teach or suggest that one or more write operations have been buffered on the host processor memory. Klaasen et al. teach away from buffering, specifically teaching "when an access operation to the disk drive is initiated, the spindle speed is increased until the disk is rotated at the normal operating velocity ... In other words, the power saving mode is disengaged prior to the commencement of read and write operations." (Col. 1, lines 51-56). It is clear that Klaasen et al. teach a system that powers down or spins down the spindles of a disk drive only until a read/write request is made, and then immediately spins up the drive. This is counter to Applicant's claimed invention which requires writing previously buffered write request. Further, the read/write requests taught by Klaasen et al. are not a *predetermined condition*, but only requests. A request is not a predetermined condition as described by Applicant. Thus, the Examiner has failed to show a *prima facie* case of obviousness and Claim 6 is believed allowable. Further, Claim 6 is allowable as being dependent on an allowable base claim.

Claim 7 is rejected under 35 U.S.C. § 103(a) as being unpatentable over Klaasen et al. in view of Olds et al. and further in view of U.S. Patent Application Publication 2003/0093645 to

09/991,164

Attorney Docket No.: P12570

Wong et al. (hereinafter, "Wong et al."). This rejection is respectfully traversed and Claim 7 is believed allowable based on the foregoing and following discussion.

As discussed above, Klaasen et al. and Olds et al. do not teach the limitations of the claim. Further, Wong et al. teach a web based system for caching large amounts of information on hard disks (Para. [0032]) and would not be properly applied to the teachings of Klaasen et al. and Olds et al. These references solve different problems and one of ordinary skill in the art would not likely combine cache techniques for physical memory on a platform to web-based disk caching of URLs and web pages for access over the Internet. Wong et al. do not solve a problem in the domain of power saving techniques for portable computers and thus is not applicable to the invention. Moreover, Wong et al. teach caching data to a disk drive and not to physical memory [0032]. Specifically Wong et al. teach "objects cached by urlfs may be stored in memory and in a disk or in a disk only." This is properly interpreted to mean that Wong et al. always store cached objects on a disk, and sometimes, additionally, store the cached objects in memory. This teaches away from Applicant's claimed invention of buffering data in physical memory, and will not result in power-savings by the computing device.

Further, none of the cited references, either alone, or in combination, teach that the request is received on a host processor and that buffering is to be performed by an intermediate FSD executing on the host processor where the predetermined condition is identified by the FSD.

The § 103 rejection of claim 7 based on Klaasen et al., Olds et al. and Wong et al. is also in error under the law. As an initial matter, the Office Action fails to identify a legally cognizable suggestion for combining Wong et al. to Klaasen et al. and Olds et al. In this regard, the Office Action states: "*it would have been obvious for one having ordinary skill in the art, having the teachings of Wong, Klaasen and Old [sic] before him at the time of the invention was made, to modify the method of Klaasen and Old [sic] to include the teachings of Wong, in order to obtain the claimed method. One of ordinary skill in the art would have been motivated to make such a combination as it provides a way to increase the efficiency of write operations to disk drives.*" However, as a matter of law and fact, this is not a proper suggestion for combining Wong et al. to Klaasen et al. and Olds et al.

09/991,164

Attorney Docket No.: P12570

Turning first to the legal error, Applicants wishes to remind the Office of the bedrock legal principles for rejecting a claim under 35 U.S.C. § 103. Specifically, in In re Rouffet, 47 U.S.P.Q.2d 1453 (Fed. Cir. 1998) the Federal Circuit explained:

To reject claims in an application under section 103, an examiner must show an un rebutted prima facie case of obviousness. In the absence of a proper prima facie case of obviousness, an applicant who complies with the other statutory requirements is entitled to a patent.

Id. at 1455 (citations omitted and emphasis added).

In the Rouffet case, the Examiner had rejected the pending claims on a combination of references. The Board sustained the Examiner. However, the Federal Circuit reversed the Board's decision and ruled that the Examiner's rejections were legally impermissible because they failed to demonstrate a suggestion for combining the references in the manner proposed by the Examiner. As explained by the Federal Circuit:

As this court has stated, "virtually all [inventions] are combinations of old elements." Therefore, an examiner may often find every element of a claimed invention in the prior art. If identification of each claimed element in the prior art were sufficient to negate patentability, very few patents would ever issue. Furthermore, rejecting patents solely by finding prior art corollaries for the claimed elements would permit an examiner to use the claimed invention itself as a blueprint for piecing together elements in the prior art to defeat the patentability of the claimed invention. Such an approach would be "an illogical and inappropriate process by which to determine patentability." To prevent the use of hindsight based on the invention to defeat patentability of the invention, this court requires the examiner to show a motivation to combine the references that create the case of obviousness.

Id. at 1457-58 (citations omitted and emphasis added). These principles have not been followed in rejecting Claim 7. Merely stating an advantage or possible advantage of combining references, as was done to reject Claim 7, is not the same as "show[ing] a motivation to combine the references."

On the contrary, in order to establish a *prima facie* case of obviousness, there must be actual evidence of a suggestion to modify a prior art reference or to combine two prior art references, and the suggestion to combine or modify the prior art must be clear and particular. In

09/991,164

Attorney Docket No.: P12570

re Dembiczak, 50 U.S.P.Q.2d 1614, 1617 (Fed. Cir. 1999). In order to establish a *prima facie* case of unpatentability, particular factual findings demonstrating the suggestion to combine must be made. See, for example, Ecolochem Inc. v. Southern California Edison, 56 U.S.P.Q.2d 1065, 1072-73 (Fed. Cir. 2000) and In re Dembiczak, 50 U.S.P.Q.2d 1614, 1617-1618 (Fed. Cir. 1999). Indeed, the law is quite clear that an obviousness rejection must be based on facts, not conjecture.

The Supreme Court... foreclosed the use of substitutes for facts in determining obviousness under section 103. The legal conclusion of obviousness *must be supported by facts*. Where the legal conclusion is not supported by facts it cannot stand.

In re Warner, 379 F.2d 1011, 1017 (C.C.P.A. 1967). This longstanding principle has been followed to date. For example, in the unpublished Board decision, Ex parte Megens, App. No. 1999-0277 (B.P.A.I. Oct. 29, 1999), the Board stated:

Rejections based on 35 U.S.C. § 103 must rest on a factual basis. In re Warner, 379 F.2d 1011, 1017, 154 USPQ 173, 177-78 (CCPA 1967). In making such a rejection, an examiner has the initial duty of supplying the requisite factual basis and may not, because of doubts that the invention is patentable, resort to speculation, unfounded assumptions or hindsight reconstruction to supply deficiencies in the factual basis. Id.

The examiner's conclusion that it would have been obvious to incline Phillips' loading dock floor 65 rests on the completely unfounded assumption that it would be desirable to drain liquid from the floor. The Phillips reference, however, is devoid of any indication that liquid might accumulate on the floor or that such accumulation would pose a problem even if it did occur. It is therefore apparent that the examiner has resorted to improper speculation and hindsight reconstruction to overcome the admitted deficiency of Phillips vis-à-vis the subject matter recited in claim 1.

(Megens at Pages 4-5)(emphasis added).

This is precisely the situation presented here. The "suggestion" in support of the rejection of Claim 7 amounts to nothing more than a speculative statement that, given the alleged presence of the claim elements in the prior art and an advantage that combining these elements would allegedly achieve, a person skilled in the art would have found it obvious to combine the references to create the claimed invention. The problem with this approach is that it effectively eliminates the requirement of identifying a suggestion for combining references from the

09/991,164
Attorney Docket No.: P12570

obviousness analysis. More specifically, the analysis present in the Office Action proceeds in the following manner:

- a) What elements are present in the pending claims?
- b) Can these elements be found in prior art references?
- c) If they can be found, and the references themselves provide no suggestion for combining these elements, can some end or advantage be identified to combine the elements in the manner proposed in the Applicants' claims?
- d) If so, combine the elements in the manner proposed by the Applicants and reject the pending claims.

This mode of analysis is, of course, deeply flawed. Specifically, as noted by the Federal Circuit in the Rouffet quote identified above, all of the elements of most claimed inventions can almost always be found in the prior art. Therefore, the answer to step "b" above will almost always be "yes". Since it is a statutory requirement that all inventions have utility, there will also always be an identifiable end or advantage in combining the elements in the prior art in the manner proposed by any claim (e.g., if there was no purpose to an element in a claim it would not be included in the claimed apparatus, after all, who would pursue a claim with superfluous elements or a claim with no utility?). Therefore, if the "suggestion" requirement of 35 U.S.C. § 103 can be met by merely identifying any end or advantage which will be achieved by combining the elements of the prior art references, the suggestion requirement can always be met and is utterly meaningless.

This inherent flaw in the analysis employed in rejecting Claim 7 is elucidated by viewing the alleged "suggestion" the Office Action identifies in support of the rejection. As noted above, in rejecting Claim 7, the Office Action states: "it would have been obvious for one having ordinary skill in the art at the time of the invention ... to modify the method of Klaasen and Old [sic] to include the teachings of Wong, in order to obtain the claimed method. One of ordinary skill in the art would have been motivated to make such a combination as it provides a way to increase the efficiency of write operations to disk drives." The first part of the statement, namely, "It would have been obvious ... to modify" is merely boilerplate language that does not address the suggestion requirement. The second part of the statement, describing the modification is omitted. This second part of the statement, is conclusory, e.g., that a mere

09/991,164
Attorney Docket No.: P12570

combination will result in the claimed invention. However, not only is the modification not described, but the motivation to combine specific elements of the references is not provided, either. The final part of the statement, namely, "one of ordinary skill in the art would have been motivated to make such a combination as it provides a way to increase the efficiency of write operations to disk drives," must, then be the alleged "motivation" for modifying combining the references.

However, while it is true that one possible advantage of buffering data may result in more efficient write operations, that is not a suggestion in and of itself for combining data buffering with Klaasen et al. and Olds et al. "The mere fact that the prior art may be modified in the manner suggested by the Examiner does not make the modification obvious unless the prior art suggested the desirability of the modification." In re Fritch, 23 U.S.P.Q.2d 1780, 1783-84 (Fed. Cir. 1992)(emphasis added). Here, the Office Action does not identify any evidence in the prior art indicating or in any way suggesting the desirability of the proposed modification. It only identifies an old element that has an inherent property. Indeed, the Office Action's naked, conclusory statement amounts to nothing more than stating "A person of ordinary skill in the art would be motivated to modify Klaasen et al. and Olds et al. to include buffering or write requests in physical memory because they would want to gain a benefit from having a buffering of write requests." In other words, the Examiner is effectively saying that the motivation of adding buffering of write requests to Klaasen et al. and Olds et al. is to have the inherent benefit of adding buffering. Of course, such circular reasoning (i.e., add "X" to have "X") cannot be a legally proper tool for identifying a suggestion for combining references. If it were, no combination of old elements would ever be patentable since one can always nakedly state, a person would be motivated to add old element X from one reference to another reference because adding element X offers an advantage (again, if adding "X" had no advantage, who would ever claim it?). Simply put, there is *always* an advantage to combining old elements that can be identified through hindsight *once that combination is known*.

It should be quite clear from the above that merely identifying an advantage for adding an old element to a combination of elements is not a proper suggestion for making that combination. The MPEP further proves this point. In particular, MPEP § 2144 states that "the strongest rationale for combining references is a recognition... in the prior art or... based on established

09/991,164
Attorney Docket No.: P12570

scientific principles or legal precedent, that some advantage would have been produced by their combination.” The MPEP cites In re Sernaker, 702 F.2d 989, 994-95 (Fed. Cir. 1983) to support this proposition.

Looking at the Sernaker case, we see that the Federal Circuit states: “The lesson of this case appears to be that prior art references in combination do not make an invention obvious unless something in the prior art references would suggest the advantage to be derived from combining their teachings.” Sernaker, 702 F.2d at 995-96 (emphasis added). Notice that this statement does not state that it is obvious to combine references simply because there is an advantage to doing so. On the contrary, it carefully states that there can be no obviousness ruling unless something in the art suggests an advantage to combining the references. The advantage itself is not the suggestion, but rather the Court makes it clear that something else suggests the advantage.

The MPEP quote noted above is similar. It states that the “strongest rationale for combining references is a recognition... in the prior art or... based on established scientific principles or legal precedent that some advantage or expected beneficial result would have been produced by their combination.” (MPEP, Page 2100-127) (emphasis added). This, of course, does not state that the strongest rationale for combining references is the mere presence of an advantage to doing so. Instead, as in Sernaker, the strongest rationale is a recognition (i.e., a suggestion) in the art that an advantage will result.

Turning back to the rejections at issue, rather than identifying something in the art that suggests an advantage to making the combination, the Office Action just looks for the advantage itself and mislabels that advantage as “suggestion.” As explained above, this is a literal elimination¹ of the suggestion requirement. Since there is always an advantage to a claimed element (or why would you claim it?), the Office Action’s misplaced view of an advantage as the suggestion inherently renders all combinations of old elements unpatentable precisely because it eliminates the suggestion requirement from the analysis. Clearly, neither the MPEP section noted above nor the Sernaker case upon which that MPEP section rests for authority stands for

¹ It literally removes the “recognition” portion of the MPEP quote above and the “something in the art” portion of the Sernaker quote.

09/991,164

Attorney Docket No.: P12570

the proposition that an advantage of an element is a suggestion in and of itself for including that element in a combination.

In view of the foregoing, applicants respectfully submit that the § 103 rejection of Claim 7 must be withdrawn because it fails to identify a legally proper suggestion for combining the prior art references in the manner proposed by the Office Action. In other words, the Office has failed to establish a *prima facie* case of obviousness under 35 U.S.C. § 103. On this basis alone, the rejections of Claim 7 and all claims depending therefrom must be withdrawn.

Claim 8 is rejected under 35 U.S.C. § 103(a) as being unpatentable over Wong et al, Klaasen et al. and Olds et al. and further in view of USPN 6,711,686 to Barrett (hereinafter, "Barrett"). This rejection is respectfully traversed and Claim 8 is believed allowable based on the foregoing and following discussion.

As discussed above, Klaasen et al., Olds et al., and Wong et al. do not teach all of the limitations of base Claim 7. Further, Wong et al. is improperly combined with the other references as there is no motivation to use techniques for web caching to result in power savings for a portable computer. Barrett teaches data stored in a disk cache (col. 2, line 7) that is flushed to the disk only upon an exit. This is contrary to Applicant's claimed invention which requires caching to physical memory. Further, Claim 8 requires receiving user input requesting that the buffered write operations be committed. A user input request to commit is not the same as receiving a system notification that the system is shutting down. Thus, the Examiner has failed to show all of the limitations of claim 8 in properly combined prior art and Claim 8 is believed allowable.

Claim 9 is rejected under 35 U.S.C. § 103(a) over Klaasen et al., Olds et al., as applied to Claim 2 and further in view of U.S. Patent Application Publication 2002/0019874 to Borr. This rejection is respectfully traversed and Claim 9 is believed allowable based on the foregoing and following discussion.

As discussed above, Klaasen et al. and Olds et al. fail to show the limitations of the base Claims. The Examiner asserts that Borr teaches *determining whether the requested write operation corresponds to an entity registered to participate in the method of controlling device write operations*. This is incorrect. Borr teaches [0107] that a file has a parameter defining its access mode, i.e., read-only, writeable, etc. The file client device, as taught by Borr, associates a

09/991.164

Attorney Docket No.: P12570

lock with the file based on whether it can be written to. This is not the same as registering to participate in a method of controlling write operations, as defined by Applicants. By definition, all files in Borr *participate* in the access control method because all have an on/off flag. In contrast, Applicant requires determining whether the entity actually *participates* in the controlling process by registering with the intermediate file system driver. It will be apparent to one of ordinary skill in the art that placing a flag in a file is not the same as proactively registering the compliance of the read/write buffering scheme with the intermediate file system driver. Applicants describe this registration scheme, at least in Para. [0017]. This registration is not a term that one of ordinary skill in the art would assume means that a flag is set in a file. Thus, Claim 9 is believed allowable.

Claim 10 is rejected under 35 U.S.C. § 103(a) over Klaasen et al., Olds et al., as applied to Claim 2 and further in view of USPN 5,815,648 to Giovannetti (hereinafter, "Giovannetti"). This rejection is respectfully traversed and Claim 10 is believed allowable based on the foregoing and following discussion.

As discussed above, Klaasen et al. and Olds et al. fail to teach the limitations of the claim. Therefore, combining the writeback operation of Giovannetti will not result in Applicant's claimed invention.

Claims 11-14, 16-20 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Klaasen et al. in view of Morcom. This rejection is respectfully traversed and Claims 11-14, and 16-20 are believed allowable based on the foregoing and following discussion.

As for Claims 11-14, the cited reference at Col. 2, lines 32-40 does not teach determining whether a limited power condition exists. Klaasen et al. teaches only that the controller reduces transfer speed when the storage device is powered by battery. This is not the same as a step of *determining* whether a limited power condition exists. Further, the Examiner cannot look at each element of the recited claim in a vacuum. The power condition determines whether only the requested data is read or whether a superset of the data is read. One cannot combine the teaching of Klaasen et al. and Morcom to provide alternate steps. Klaasen et al. specifically teach that the spindle speed is to operate at reduced rotation and/or a lower transfer rate when operating under reduced power. At no time is it suggested that that varying the amount of data to be read could be an option. Because the purpose of Klaasen et al. is to reduce the power consumption of the

09/991,164

Attorney Docket No.: P12570

device by reducing the spindle rotation, applying the teaching of Morcom (to copy a maximum amount of data to cache) is counter-intuitive and would not result in Applicant's invention. Applying Morcom to Klaasen et al. would increase the power consumption by requiring excessive data to be read to fill the cache, and thus not fall within the teaching and intent of Klaasen et al. Further, Morcom teaches copying a maximum amount of data to a cache and then immediately powering down the disk regardless of the power state of the device before the read. This is counter to Applicant's claimed invention which does not immediately power down the device regardless of the power state. Thus, Morcom is misapplied to Klaasen et al. and a combination of the two references will not result in Applicant's invention.

In addition, Applicant's claim requires that the file portion is to be read into memory. The gist of the invention is to optimally access the storage device and avoid activating the device when unnecessary. Morcom teaches reading data to cache memory and not directly into physical, or system memory so that the cache is always full. This is counter to the claimed invention and also counter to the purpose of Applicant's invention. Applying Morcom would require that the cache be filled with data regardless of whether the data is logically related to the requested portion. This relationship is what drives whether a superset, a subset or entire file are read. Moreover, neither Klaasen et al. or Morcom teach an intermediate file system driver as recited in the amended claim, and as discussed above.

Specifically with regard to Claim 14, the Examiner attempts to introduce set theory and assert that a "subset" of the file is the same as the entire file. In light of the USPTO requirement of claim differentiation, the Examiner's interpretation cannot be the case. Claim 12 specifically requires the superset to be the entire file. Thus, as would be apparent to one of ordinary skill in the art, the subset will be understood to be a portion smaller than the entire file. Therefore this rejection is improper. Thus, Claims 11-14 are believed allowable.

As for Claim 16, Morcom teaches that the data is returned from cache memory. In contrast, Applicant's claim requires that the data be read from the device and not cache memory. Thus, applying Morcom would return data retrieved from cache memory to the process and would preclude returning data retrieved from the device. The Examiner asserts that data retrieved from cache memory is data retrieved from the device. The Examiner seems to assume that the cache memory is located on the device. However, base Claim 11 requires the buffer

09/991,164

Attorney Docket No.: P12570

memory to be coupled to the processor. Thus, cache memory located on the device is not the same as the memory recited in the Claim. Thus, when the claim requires accessing the device to read the requested file portion and then returning the data to the process, there is no reading of data from the cache memory, but merely data retrieved from the device.

As for Claim 17, Morcom does not teach that the requested portion is returned to the requesting process before a remainder is read. As discussed above, Morcom does not teach a superset as defined by Applicant, and Morcom does not teach the temporal requirements of returning the requested portion first. Morcom merely teaches reading as much data as will fit in the cache. Further, the Examiner's response that it is inherent in Applicant's limitation that a superset might be read in some cases is not relevant. The limitation that is not taught or suggested by Morcom is selectively reading a superset by the intermediate FSD. Morcom reads data until the cache is full. There is no selection determination at all.

Claim 18 is allowable as being dependent on an allowable base claim, as discussed above.

As for Claim 19, Morcom does not teach or suggest *if a superset of the requested file portion is read into memory, deactivating the device*. Morcom teaches that cache is always filled to the maximum. In some cases, this might even be less than a requested portion. Further, Morcom teaches that the device is always deactivated after a read. Morcom does not teach an inherent determination that a superset of the requested portion has been read into memory or selectively choosing to read a superset by the intermediate FSD. Nor does Morcom teach that there is a possibility that the device will not be inactivated, which is inherent in Applicant's claimed invention. Moreover, Morcom teaches only that data be read into a cache and not that a requested file portion be read into memory.

Claim 20 is allowable as being dependent on an allowable base claim, as discussed above. Thus, Claims 11-14 and 16-20 are believed allowable.

Claims 15, 22-23 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Klaasen et al. and Morcom and further in view of U.S. Patent Application Publication 2002/0091902 to Hirofuji (hereinafter, "Hirofuji"). This rejection is respectfully traversed and Claims 15, 22-23 are believed allowable based on the foregoing and following discussion.

Klaasen et al. and Morcom do not teach the recited limitations, as discussed above.

09/991.164

Attorney Docket No.: P12570

As for Claim 15, Hirofuji does not teach or suggest identifying the subset of the file to be read into memory is based on one or more file access trends. Hirofuji actually discusses the arrangement of the data in the cache memory. Further, Hirofuji discusses the reading of information into cache memory in an arrangement based on likely access order, i.e., sequential or random. As described in the Specification, at least in para. [0037], Applicant describes temporal trends: i.e., which files, and which portions of files, have been accessed recently. Hirofuji does not teach or suggest temporal trends, but merely characteristics of a file. Morcom teaches to read data until cache memory is full. Even if it could be surmised that Hirofuji teaches selection of data based on file access trends, it is improper to combine this with Morcom. It is improper to combine the teaching of Morcom with Hirofuji as they teach incompatible methods; thus, there is no motivation to combine, as discussed above. One cannot select a portion of a file and also fill cache until it is full. The Examiner asserts that it is possible to fill a cache with a portion of a file. However, it is not possible to do this every time, which is what would result by combining the teachings of Morcom and Hirofuji. One concept is to reduce and one concept is to increase. This result is contrary to Applicant's claimed invention. Thus, the rejection is improper and should be withdrawn.

As for Claims 22-23, Hirofuji does not teach the registering of file types with the intermediate FSD, as described by Applicant. This term is specifically discussed in the Specification in the context of Applicant's invention and other definitions cannot be assumed by one of ordinary skill in the art.

As for Claim 23, Hirofuji teaches an arrangement priority to distinguish between likely sequential or random access. This is not the same as a relative file type priority and selectively storing a superset file portion relative to the relative priority. Hirofuji teaches that data may be stored differently based on its access type. Applicant requires that a superset of the requested file portion may be selectively stored based on a relative priority. Thus, the Examiner has failed to show a *prima facie* case of obviousness and Claims 15 and 22-23 are believed allowable.

Claims 24-26 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Ryu in view of Klaasen et al. This rejection is respectfully traversed and Claims 24-26 are believed allowable based on the foregoing and following discussion.

09/991,164

Attorney Docket No.: P12570

As discussed above, Ryu teaches the automatically storing of data being used and to turn off the system in order to prevent the loss of data, i.e. hibernation. Ryu teaches detecting whether battery voltage is between a specific voltage range. Ryu teaches *selecting* a power state, i.e., power down when battery voltage is dropping. Ryu does not teach the step of *detecting* that a time-out period is about to expire and that the system is operating under a limited power condition. Ryu's device may be operating under battery power (a limited power condition as defined by Applicant), but not detect a time-out, as required in Claims 24-26. Further, Ryu teaches always storing – not buffering – data when the system battery is low. The prevention of data loss is not the same problem as optimizing power consumption and setting the device to deactivate at a specified time-out period. While Klaasen et al. teach that a spindle may be deactivated after a predetermined amount of time, there is no motivation to combine these references, as required by law, as discussed above. Ryu teaches a system to save data before a system is shutdown due to lack of battery power. Klaasen et al. teach a system to deactivate a spindle in a device to reduce power consumption. Neither reference suggests that its teachings could be applied to be combined with the alternate purpose. Further, neither reference teaches, *wherein the writing is in response to an intermediate file system driver detecting that a time-out has occurred*. Thus, Claim 24 and its progeny are believed allowable.

As for Claim 26, as discussed above, Ryu does not teach determining whether a limited power condition exists, but assumes battery power and just compares voltages to determine when the battery voltage is low enough to force a hibernation mode. Thus, Claims 24-26 are believed allowable.

Claim 21 is rejected under 35 U.S.C. § 103(a) as being unpatentable over Klaasen et al. and Morcom in view of USPN 6,442,647 to Morton et al. (hereinafter, "Morton et al."). This rejection is respectfully traversed and Claim 21 is believed allowable based on the foregoing and following discussion.

Klaasen et al. do not disclose each and every limitation as asserted by the Examiner, as discussed above. Further, Morton et al. teach a system to take a request for data and turn it into two commands. The first request relates to the actual data request and the second relates to reading the remainder of data on the track. Morton et al. do not teach *wherein the superset of the requested file portion is logically related to the requested portion*. Morton et al. teach only that

09/991,164

Attorney Docket No.: P12570

additional data on the track is read. Often this data will not be related to the requested data and especially not *logically related*. Further, the Examiner has failed to address this difference in the current Office Action. The method taught by Morton et al. is brute force to maximize a data read, but does not optimize the device access. By reading only data that is logically related to the data requested, access to the device is optimized, i.e., not over-accessed for an excessive amount of data (remainder of track). Thus, Morton et al. do not teach or suggest the recited claim limitations and Claim 21 is believed allowable.

Claim 33 is rejected under 35 U.S.C. § 103(a) as being unpatentable over Morcom in view of USPN 6,415,359 to Kimura et al. (hereinafter, "Kimura et al."). This rejection is respectfully traversed and Claim 33 is believed allowable based on the foregoing and following discussion.

Contrary to the Examiner's assertion, Morcom does not teach or suggest all of the elements of the Claimed invention, as discussed in relation with Claim 31, above. Specifically, Applicant's claimed invention determines the current power state of the device and then based on that state, selectively reading a superset of the requested file portion into physical memory. Morcom reads data until the cache is full, but there is no requirement that the data is related to the requested portion, nor does Morcom teach that a lesser amount of data be read. This teaches away from selectively reading a superset. Therefore, regardless of the teaching of Kimura et al., the Examiner has failed to put forth a prima facie case of obviousness, as all of the limitations of the claim have not been shown by the cited references. Therefore, Claim 33 is believed allowable.

Claim 34 is rejected under 35 U.S.C. § 103(a) as being unpatentable over Morcom in view of Morton et al. This rejection is respectfully traversed and Claim 34 is believed allowable based on the foregoing and following discussion.

Morcom does not teach the limitation of the recited Claim, as discussed above for Claim 31. Further, Morton et al. teach a system to take a request for data and turn it into two commands. The first request relates to the actual data request and the second relates to reading the remainder of data on the track. Morton et al. do not teach *wherein the superset of the requested file portion is logically related to the requested portion*. Morton et al. teach only that additional data on the track is read, as discussed above, at least for Claim 21. Therefore, the

09/991,164

Attorney Docket No.: P12570

Examiner has failed to put forth a *prima facie* case of obviousness, as all of the limitations of the claim have not been shown by the cited references. Thus, Claim 34 is believed allowable.

Claim 41 is rejected under 35 U.S.C. § 103(a) as being unpatentable over Rao in view of Giovannetti. This rejection is respectfully traversed and Claim 41 is believed allowable based on the foregoing and following discussion.

Rao does not teach the limitations of the Claim, at least as discussed above in conjunction with Claim 39. At least, Rao teaches away from Applicant's claimed invention. Rao does not teach or suggest selectively buffering write requests to physical memory. Rao teaches always buffering disk writes to buffers within a SCSI controller board, not selectively buffering write request until a predetermined condition is detected. Further, Giovannetti does not teach that write requests are intercepted and buffered by the intermediate file system driver. Therefore, a combination of the cited references will not result in Applicant's claimed invention and Claim 41 is believed allowable.

All claims remaining in the application are now allowable.

RECEIVED
CENTRAL FAX CENTER

08:54:08 p.m.

01-03-2007

54 /54

JAN 03 2007

09/991,164

Attorney Docket No.: P12570

CONCLUSION

In view of the foregoing, Claims 1-37, 39 and 41-42 are all in condition for allowance. If the Examiner has any questions, the Examiner is invited to contact the undersigned at (703) 633-6845. Early issuance of Notice of Allowance is respectfully requested. Please charge any shortage of fees in connection with the filing of this paper, including extension of time fees, to Deposit Account 50-0221 and please credit any excess fees to such account.

Respectfully submitted,

Dated: 3 Jan. 2007/ Joni D. Stutman-Horn/

Joni D. Stutman-Horn, Reg. No. 42,173

Patent Attorney

Intel Corporation

(703) 633-6845

Intel Corporation
4040 Lafayette Center Drive
MS LF2
Chantilly, VA 20151

BEST AVAILABLE COPY

1/4




FIG. 1

2/4

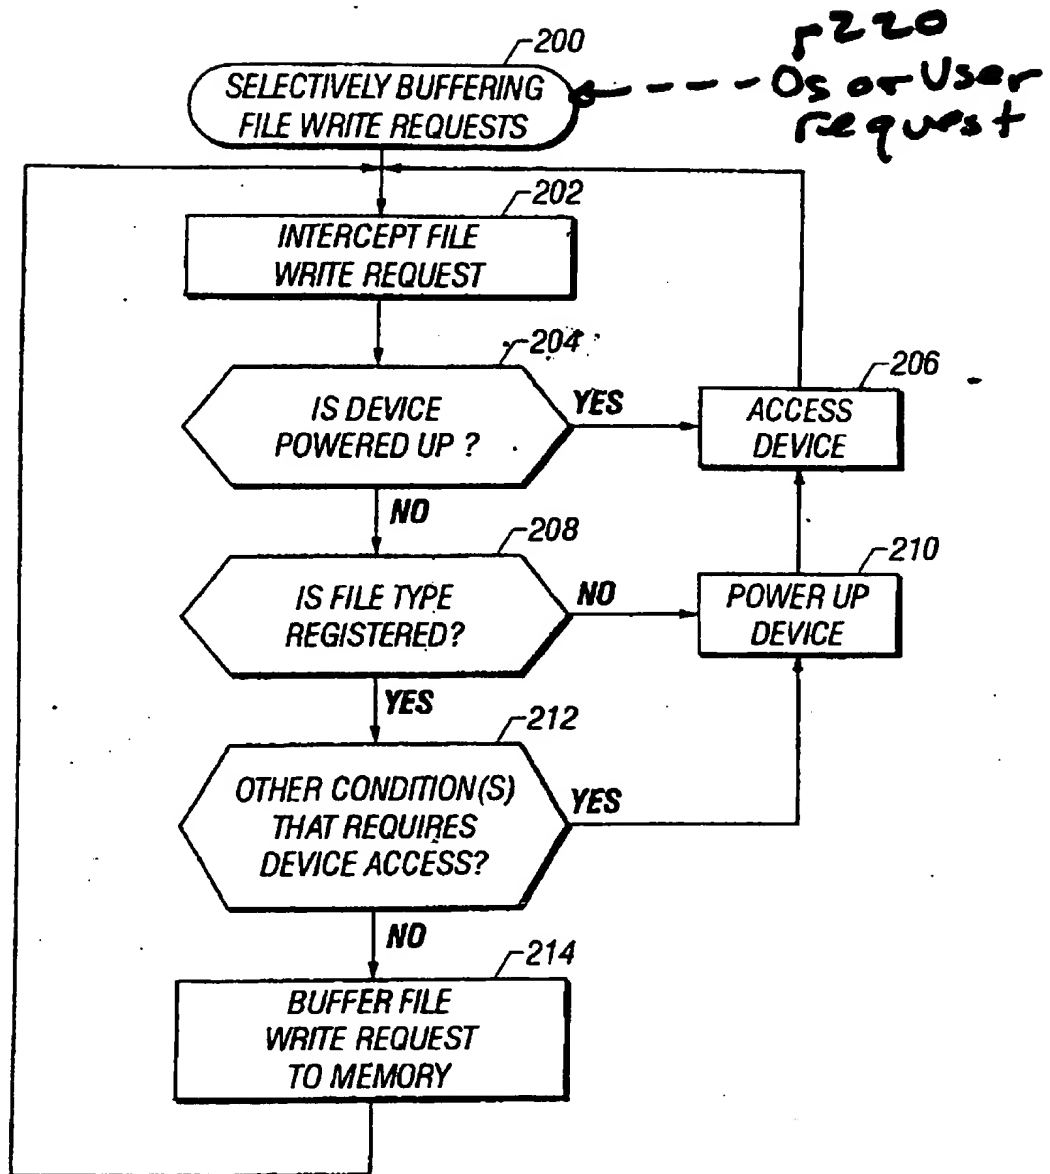


FIG. 2

BEST AVAILABLE COPY

